XT Node Architecture

Let's Review: Dual Core v. Quad Core

Dual Core

- Core
  - 2.6Ghz clock frequency
  - SSE SIMD FPU (2flops/cycle = 5.2GF peak)
- Cache Hierarchy
  - L1 Dcache/Icache: 64k/core
  - L2 D/I cache: 1M/core
  - SW Prefetch and loads to L1
  - Evictions and HW prefetch to L2
- Memory
  - Dual Channel DDR2
  - 10GB/s peak @ 667MHz
  - 8GB/s nominal STREAMs

Quad Core

- Core
  - 2.1Ghz clock frequency
  - SSE SIMD FPU (4flops/cycle = 8.4GF peak)
- Cache Hierarchy
  - L1 Dcache/Icache: 64k/core
  - L2 D/I cache: 512 KB/core
  - L3 Shared cache 2MB/Socket
  - SW Prefetch and loads to L1,L2,L3
  - Evictions and HW prefetch to L1,L2,L3
- Memory
  - Dual Channel DDR2
  - 12GB/s peak @ 800MHz
  - 10GB/s nominal STREAMs

Cray XT4 Node

- 4-way SMP
- >35 Gflops per node
- Up to 8 GB per node
- OpenMP Support within socket

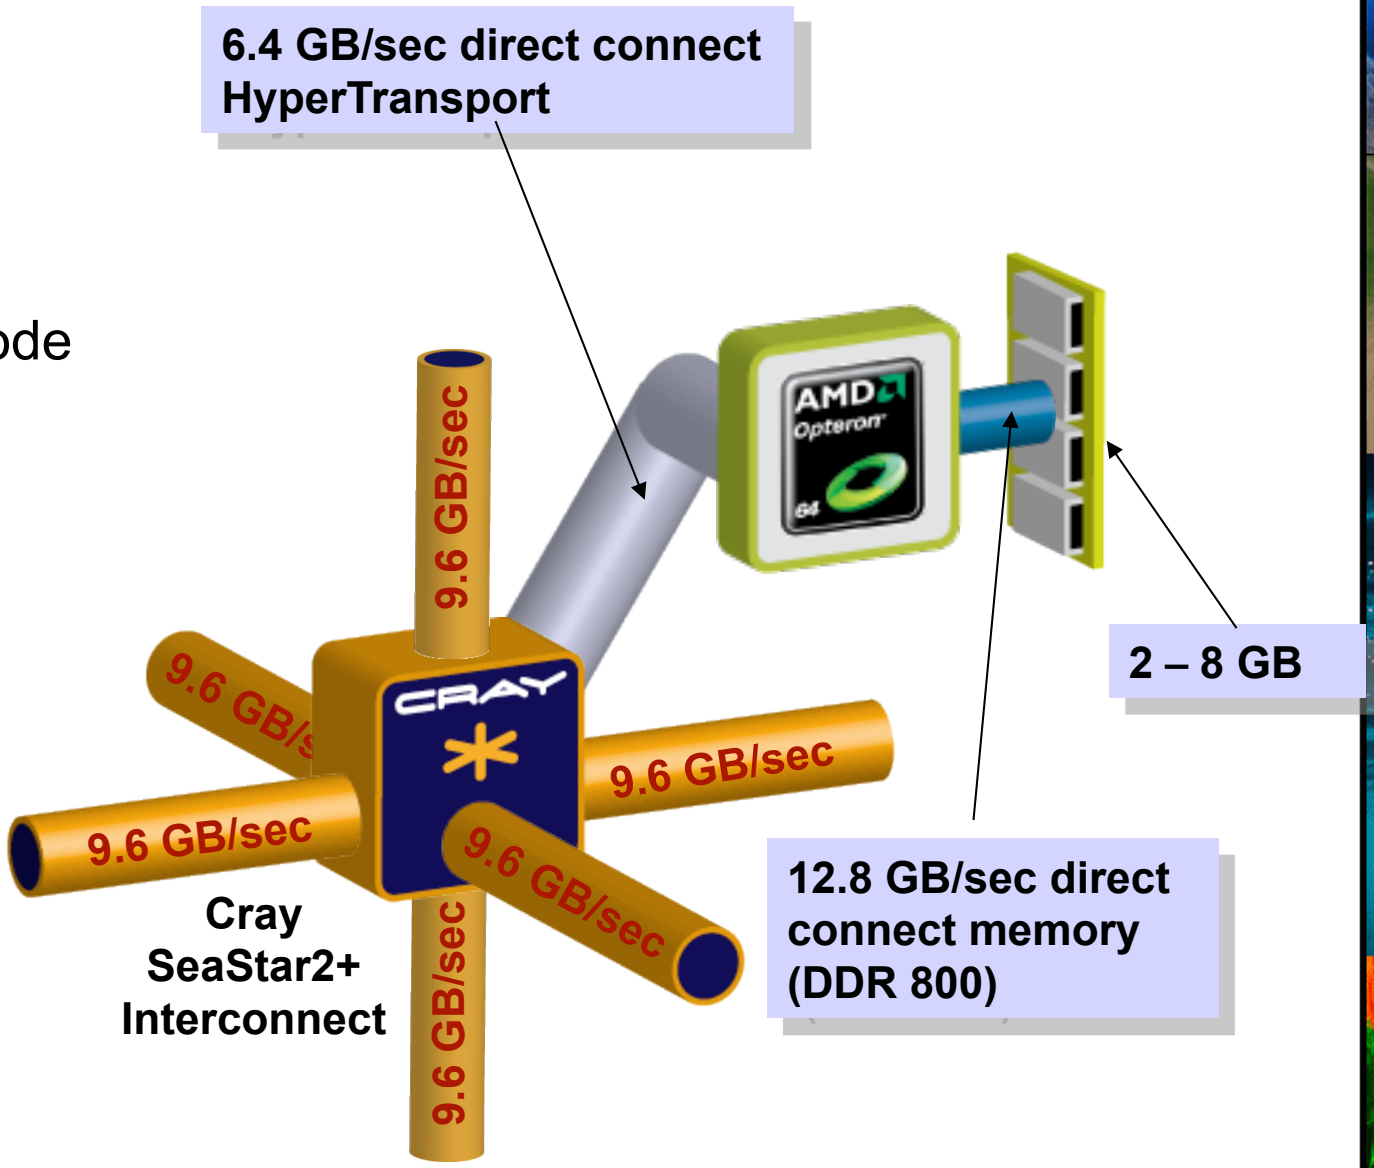

Cray XT5 Node

- 8-way SMP
- >70 Gflops per node
- Up to 32 GB of shared memory per node
- OpenMP Support

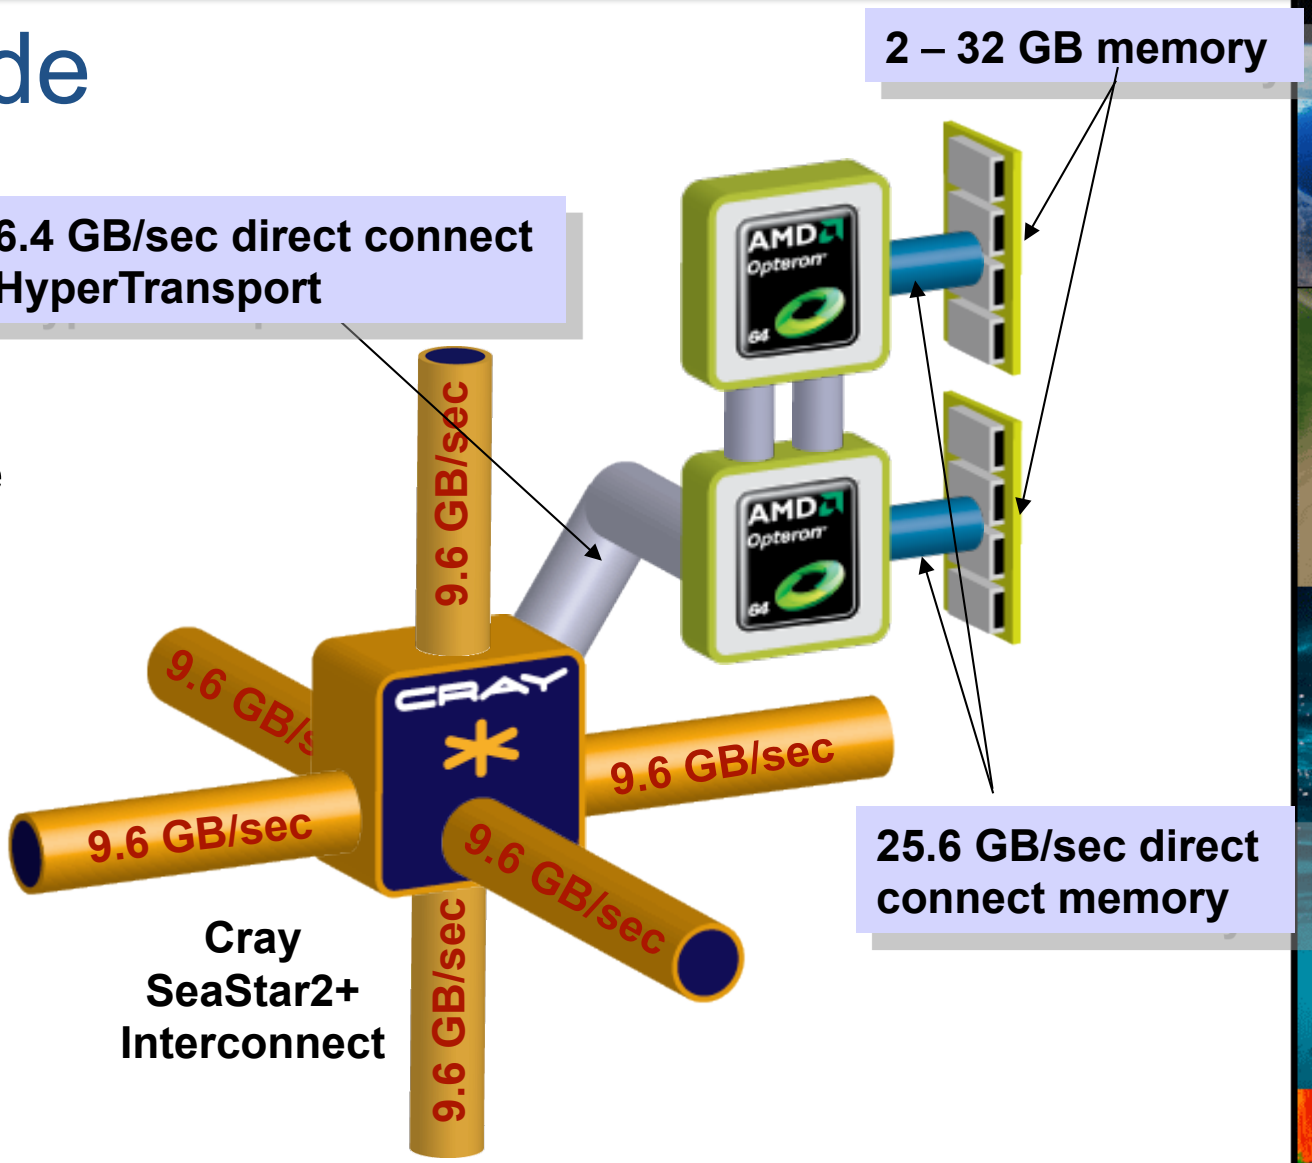

AMD Opteron Processor –Dual Core

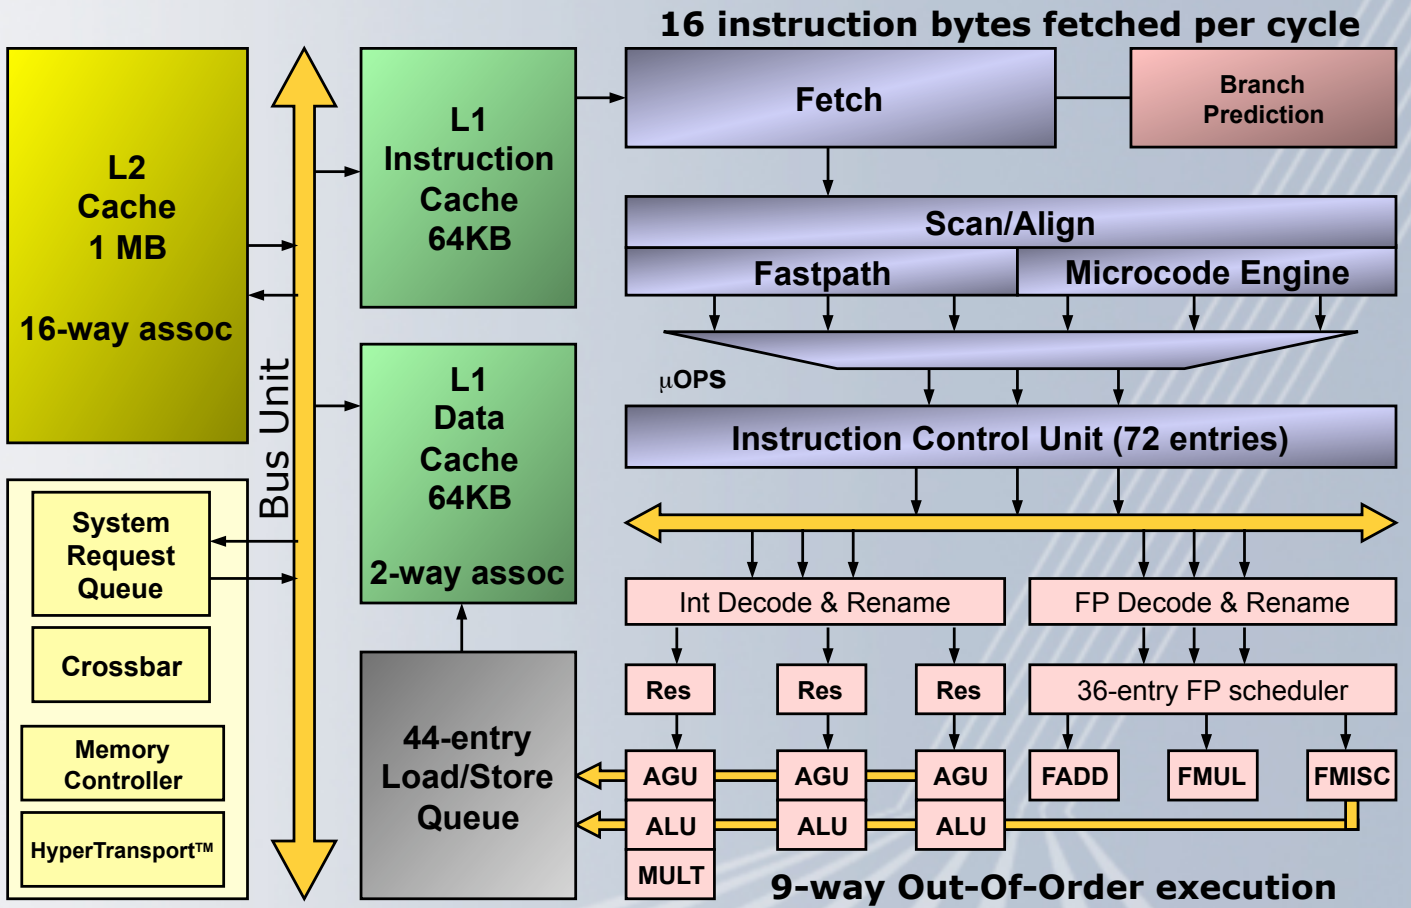

- 36 entry FPU instruction scheduler
- 64-bit/80-bit FP Realized throughput (1 Mul + 1 Add)/cycle: 1.9 FLOPs/cycle
- 32-bit FP Realized throughput (2 Mul + 2 Add)/cycle: 3.4+ FLOPs/cycle

Simplified memory hierachy on the AMD Opteron – Dual Core

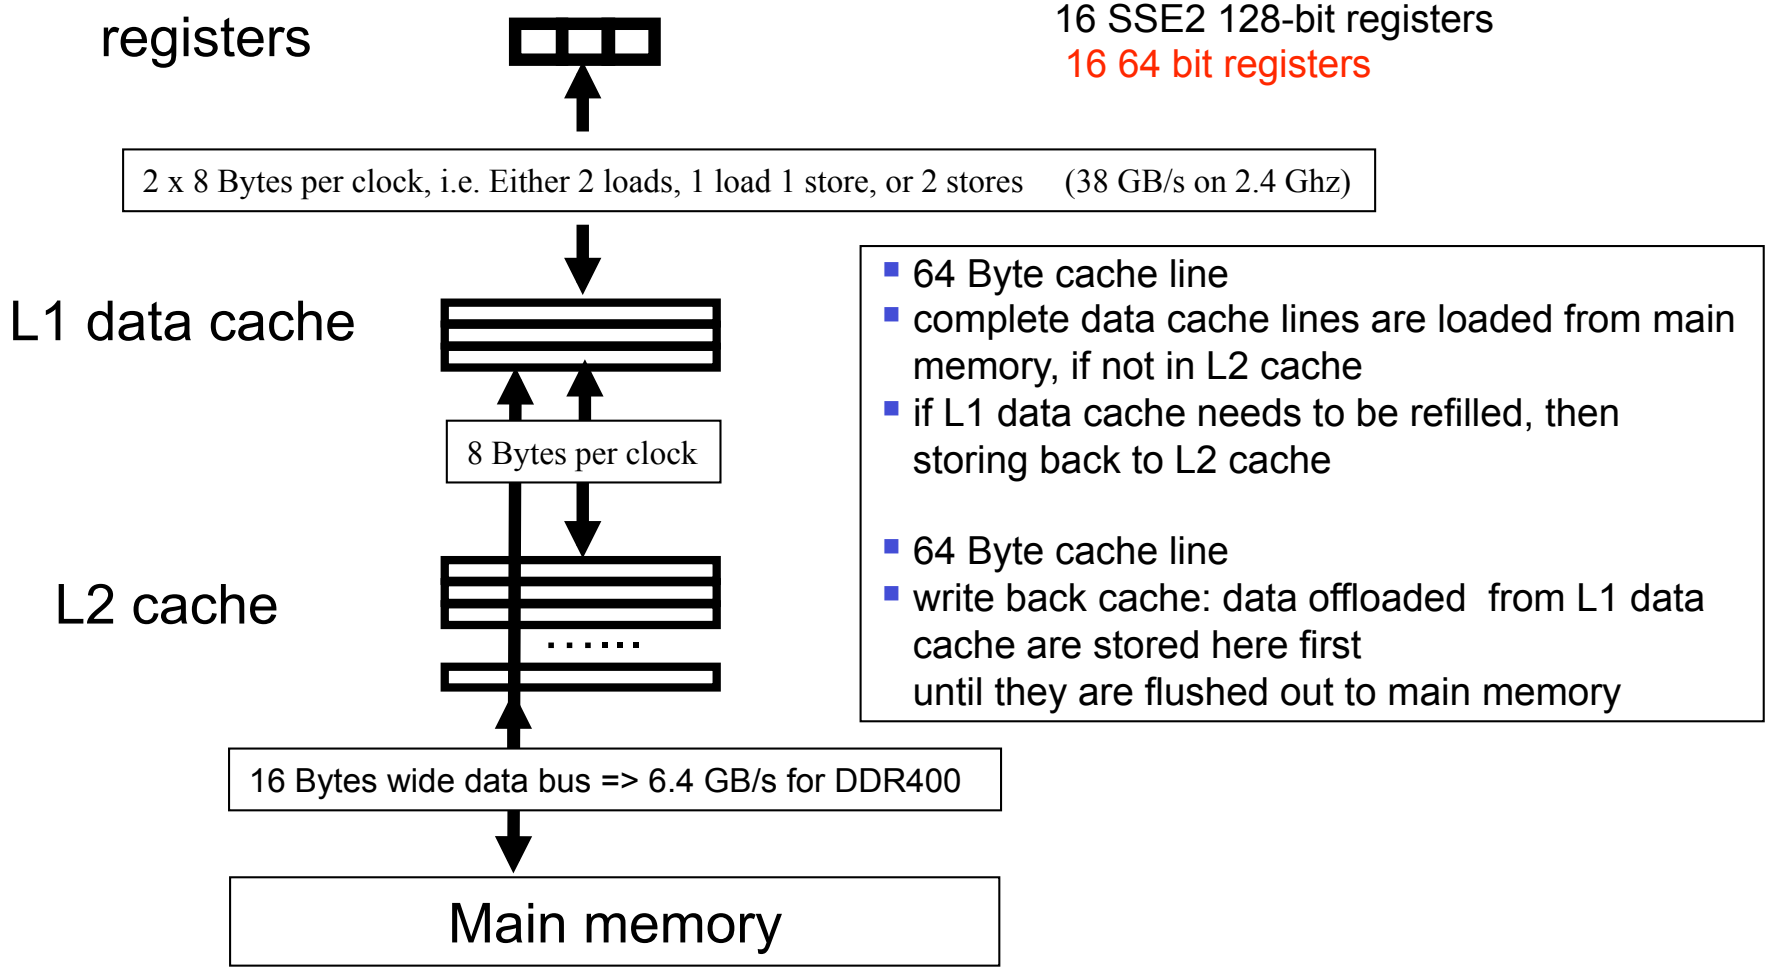

- 64 Byte cache line
- complete data cache lines are loaded from main memory, if not in L2 cache
- if L1 data cache needs to be refilled, then storing back to L2 cache

- 64 Byte cache line
- write back cache: data offloaded from L1 data cache are stored here first until they are flushed out to main memory

Core IPC Improvements – Quad Core

- Improve Branch Prediction.
- TLB enhancements.
- More out of order Ld/St capability.
- New Instructions
  - POPCNT / LZCNT
  - EXTRQ / INSERTQ
  - MOVNTSD / MOVNTSS
- Fastpath support for FP to Integer data movement.

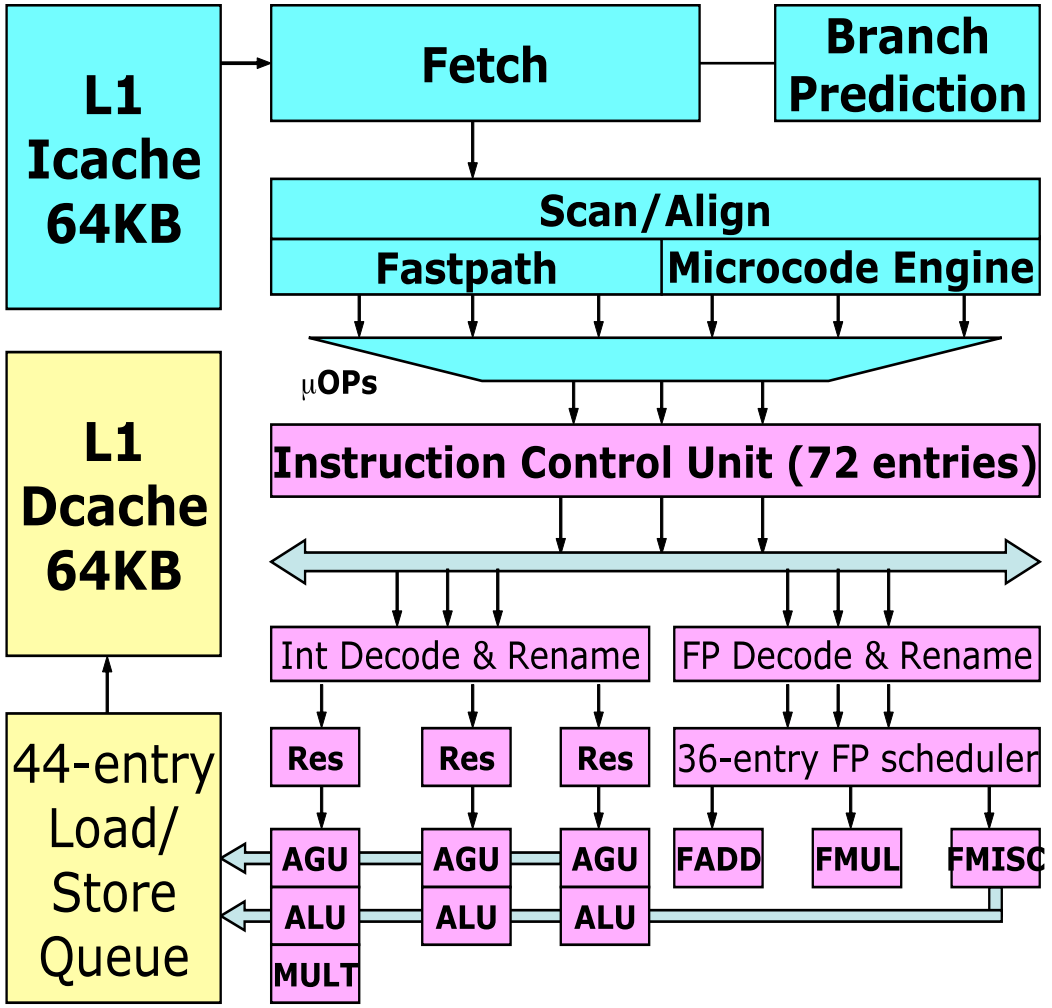

Core IPC - FastPath ? Macro-Ops? Micro-Ops?

What are these? Do I care? – (how to talk to the compiler guy)

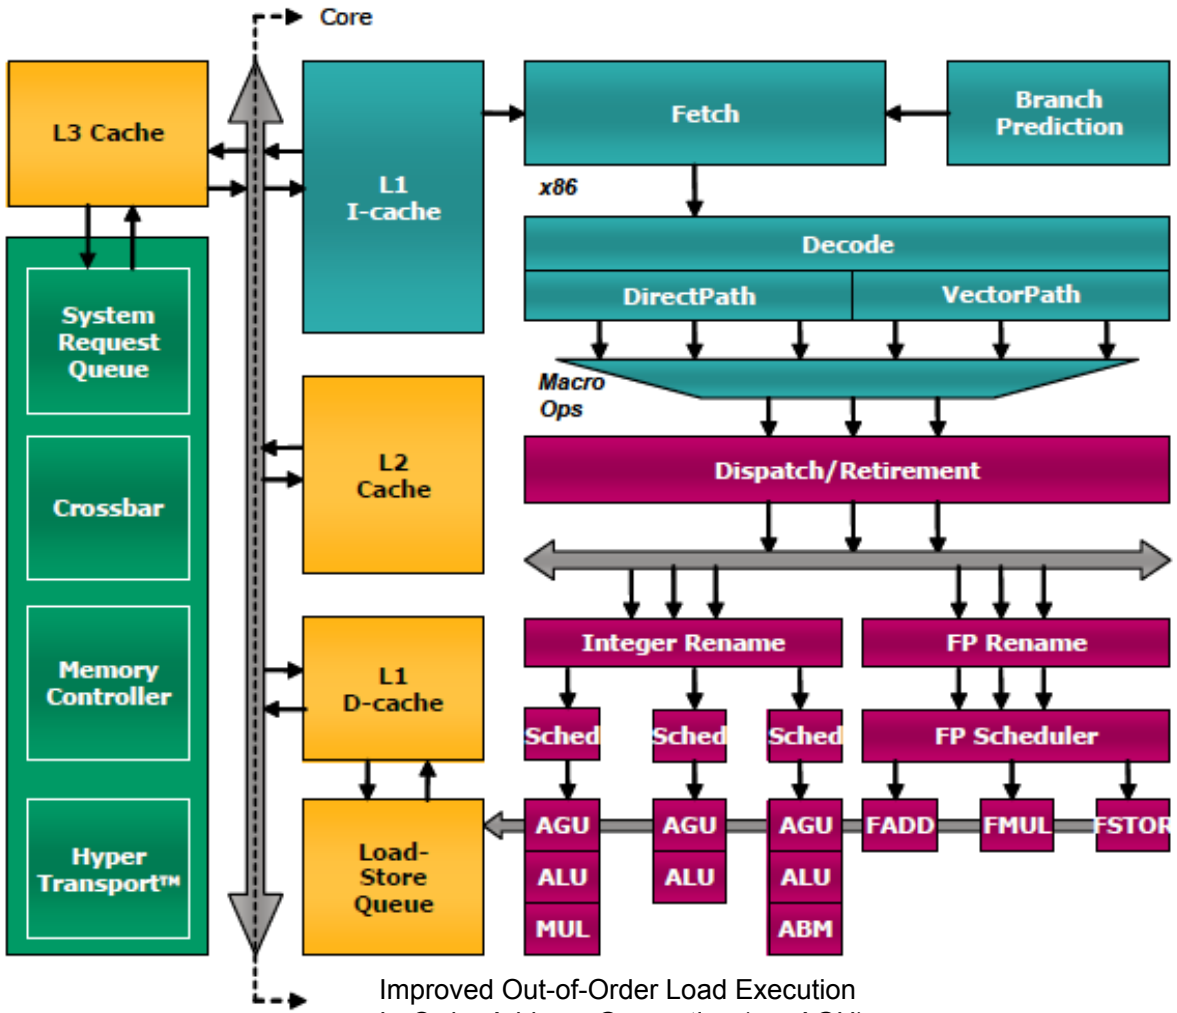

Quad Core

X86 Insts.

Avoid having more than 2 or 3 branches per 16B of instructions.

Three Decode Categories (FastPath also called DirectPath)
- DirectPath Single - *best*
- DirectPath Double - *better*
- VectorPath (microcode) - *good*

Macro-Ops

Macro-Ops tracked ReOrder Buffer (ROB).
- 72 entries (3 wide x 24 deep)
- In-order dispatch, retirement

Micro-Ops

Micro-Ops issue from Sched to Execution Units
- "Sched" aka "Reserv. Station"
- Out-of-order issue
- FP scheduler shared across units
- INT Schedulers are "per unit"

Improved Out-of-Order Load Execution
In-Order Address Generation (per AGU)
Store-to-Load Forwarding Support

Simplified memory hierachy on the Quad Core AMD Opteron – Quad Core

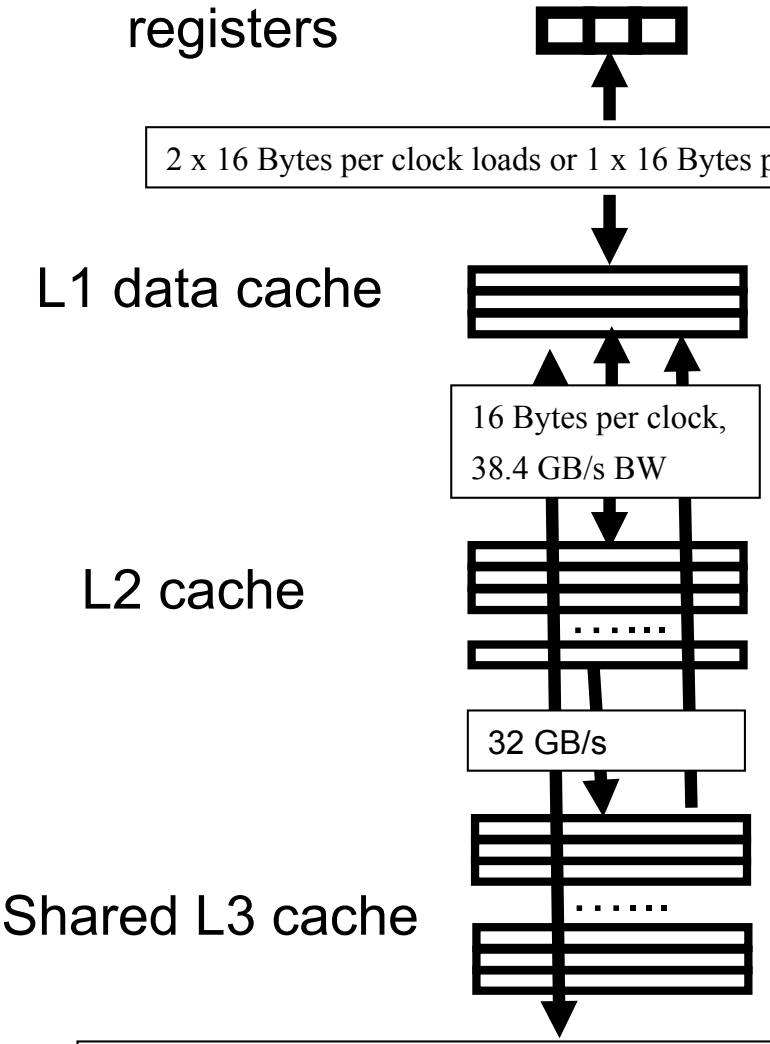

- 64 Byte cache line
- complete data cache lines are loaded from main memory, if not in L2 or L3 cache
- if L1 data cache needs to be refilled, then storing back to L2 cache, if L2 needs to be refilled, storing back to L3
- Items in L1 and L2 are exclusive, L3 is “sharing aware”

- 64 Byte cache line
- write back cache: data offloaded from L1 data cache are stored here first until they are flushed out to main memory

BACK TO TUTORIAL